

REMARKS

Claims 1-19 are pending. Claims 1-19 stand rejected.

35 U.S.C. 102 Rejections

Claims 1 and 6 stand rejected under 35 U.S.C. 102(e) as being anticipated by Hoof (U.S. Pat. No. 7,203,193, hereinafter “Hoof”). Applicants respectfully disagree.

Before addressing the Office Action rejections, Applicants believe a brief overview of Hoof is worthwhile. Hoof is directed to synchronizing notification messages to allow an even distribution of the notification messages (Hoof, abstract). Hoof Fig. 1 shows ingress control units (ICUs), a switch fabric including packet buffer units (PBUs), and egress control units. Upon receipt of a packet by an ICU, the ICU forwards the packet to a PBU for storage (Hoof col. 3, lines 57-58). The PBU stores the packet in memory and transmits a notification to all ECUs that may be interested in receiving the packet (col. 3, lines 58-60). If a particular ECU wants to receive the packet, that ECU transmits an indication (called a “booking message”) to the PBU that causes the PBU to keep the packet in memory until it is requested (col. 3, lines 64-67). Later, ECUs that want to receive the packet transmit a request to the PBU to retrieve the packet (col. 4, lines 1-3). The PBU then retrieves the packet and transmits it to the requesting ECUs for forwarding via egress ports (col. 4, lines 3-5).

Thus, Hoof uses a subscription/notification model in which incoming packets are stored and only retrieved if needed, based on subscriber ECUs that are notified of the packets, rather than pushed to egress ports as soon as possible. Hoof is directed to spreading out notifications of incoming packets in time to achieve a constant, non-bursty load on receiving (subscriber) ECUs (col. 7, lines 12-14). Hoof uses notification queues to transmit notifications in this constant, non-bursty manner (col. 7, lines 27-34).

Hoof does not disclose a “switch fabric storing cells [moved from input queues to the switch fabric] based on ... output queues,” as required by Applicants’ Claim 1 (emphasis added). The Office Action states, at page 2, part 4, that Hoof’s switch fabric 14 is equivalent to the switch fabric of Applicants’ Claim 1 because “[Hoof’s switch fabric 14] multiplexes packet data

between input and output queues.” However, as discussed above, upon receipt of a packet by an ICU, the ICU forwards the packet to a PBU (part of the switch fabric 14), where the packet is immediately stored (Hoof col. 3, lines 57-59); such storage is not based on output queues. The Office Action has equated Hoof’s egress ports 22 (e.g., as in Hoof Fig. 4) with the “output queues” of Applicants’ Claim 1, but Hoof’s egress ports do not come into the picture yet at the time of storage at Hoof’s switch fabric 14. Rather, Hoof stores each received packet in memory at a PBU and subsequently transmits a notification to ECUs that may be interested in receiving the packet (Hoof col. 5, lines 22-28). Storage of packets at Hoof’s PBU is not (and cannot be) based on Hoof’s ECUs, much less output queues at ECUs (i.e., Hoof’s egress ports 22, according to the Office Action’s claim construction).

Furthermore, due to the subscription/notification model used by Hoof, one of ordinary skill in the art would not have been motivated to modify Hoof to make storage of packets at a PBU based on Hoof’s egress ports. Hoof’s PBUs do not care about the status of output queues at ECUs, because the role of the PBUs is to store packets. Later, if an ECU that wants to receive a packet makes such a desire known via a request message, the packet is retrieved by the PBU and transmitted to the ECU for forwarding via egress ports 22. However, the egress ports only come into the picture at this last stage, and changing Hoof to make storage at a PBU based on egress ports would require substantial modification (and would not serve any purpose). The very point of a subscription/notification model as in Hoof is to keep storage at PBUs separate from functionality at ECUs, since only ECUs that subscribed to a certain message should even be involved in the process (after notification). Involving ECUs (and egress ports) at the storage stage would defeat the point of a subscription model/notification, i.e., would render Hoof unsuitable for its intended purpose, since doing so would presuppose that those ECUs care about the messages (and they might not, in reality). Such a presupposition would eliminate the efficiency gains obtained via Hoof’s subscription/notification model, which only performs processing associated with stored packets on an as-needed basis.

Since Hoof does not disclose all the elements of Claim 1, Applicants respectfully submit that the rejection of that claim under 35 U.S.C. 102(e) is improper and should be withdrawn.

Dependent Claim 6 inherits the foregoing patentably distinguishing elements of Claim 1 and should be allowed for at least the same reasons presented above.

35 U.S.C. 103 Rejections

Claims 2, 3, 8, 9, 11, 12, and 17-19 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Hoof in view of Chao et al. ("Delay-Bound Guarantee in Combined Input-Output Buffered Switches," hereinafter "Chao"). Applicants respectfully disagree.

Regarding Claims 2 and 3, these dependent claims should be allowed at least for inheriting the foregoing patentably distinguishing elements ("switch fabric storing cells... based on ... output queues") of Claim 1, since Chao does not cure the above-mentioned deficiency of Hoof for the following reasons. Chao teaches determining a virtual finishing time for packets, wherein the virtual finishing time of a packet determines when a cell is moved from an input port, straight through a switch, to an output buffer (Chao pages 517 and 519-521 and Figs. 1 and 3). Chao does not even discuss memory of a switch, much less storage of packets in memory of a switch; in fact, the Office Action of March 20, 2007 had to assert that Chao's switch inherently has memory. The foregoing arguments were previously presented in Applicants' Amendment dated June 20, 2007. Since Chao passes packets straight through from an input port to an output buffer, without storage at a memory of a switch, Chao cannot disclose a "switch fabric storing cells ... based on ...output queues," as in Applicants' Claim 1.

Claim 2 is further patentable at least for reciting the following elements:

[a] controller orders cells stored in said switch fabric based on times of said cells to depart, wherein cells having lower times of said cells to depart have higher output priorities.

As discussed above, Chao does not stores cells at Chao's switch fabric but rather passes them straight through to an output buffer. Therefore, Chao cannot teach storage at a switch fabric based on departure times.

Furthermore, the Office Action of July 10, 2008 conceded, at page 3, lines 3-4, that Chao does not disclose "output priorities for cells from [the] switch fabric to output queues." Since Chao does not teach such output priorities, Chao cannot teach the elements of Claim 2 relating to such output priorities. Furthermore, one of ordinary skill in the art would not have been motivated to combine Chao with Hoof anyway, since Hoof is directed to a technique for storing packets at a switch memory, in accordance with a subscription/notification model, in which packets are only retrieved as requested by subscribers, whereas Chao is directed to direct

transmission of packets across a switch from an input port to an output buffer, without intermediate storage at the switch fabric. The two approaches are very different and are both different from Applicants' approach, since Applicants' approach does not involve subscription/notification and does storage cells at a switch fabric. In fact, Chao teaches away from Applicants' claimed invention, since Chao teaches that performance gains can be achieved without storing cells in a switch fabric, which would have motivated one of ordinary skill in the art not to seek an architecture employing such storage.

Claim 8 recites the same patentably distinguishing elements as Claim 1 regarding storage of cells based on output queues. Chao does not cure the deficiency of Hoof regarding Claim 8, and it would not have been obvious to combine Chao with Hoof anyway, for at least the same reasons presented above. Therefore, Applicants respectfully submit that the rejection of Claim 8 under 35 U.S.C. 103(a) is improper and should be withdrawn.

Dependent Claims 9, 11, and 12 inherit the foregoing patentably distinguishing elements from base Claim 8 and should be allowed for at least the same reasons presented above.

Furthermore, Claim 9 is additionally patentable at least for reciting "updating cells after highest priority cells in said switch fabric are transferred from said switch fabric to said output queues." Chao cannot teach or suggest these elements because, as discussed above, a previous Office Action conceded that Chao does not teach "output priorities for cells from [the] switch fabric to output queues."

Regarding Claim 17, the Office Action has not presented a proper rejection of that claim. Rather, at page 7, lines 10 through 20, the Office Action has simply restated the rejection of Claim 2. However, Claim 17 is very different from Claim 2. The elements of Claim 17 are not addressed at all in the Office Action, on page 7 or elsewhere.

However, Applicants will still argue below why Claim 17 is patentable. Claim 17 recites in part:

selecting and transferring highest priority cells stored in memory
of a switch fabric from said switch fabric to output queues ...
selecting and transferring highest priority cells stored in input
queues from said input queues to said memory of said switch fabric based
on said flow-control information ...

The foregoing elements of Claim 17 relate to storage of cells in memory of a switch fabric and priorities of cells transferred from a switch fabric to output queues (output priorities). As discussed above, Chao does not store cells in memory of a switch fabric and does not teach “output priorities for cells from [the] switch fabric to output queues” (conceded in a previous Office Action). Furthermore, it would not have been obvious to combine Chao with Hoof due to the different architectures (switch that passes packets straight through to an output buffer vs. intermediate storage at PBU) and purposes (pass data through a non-blocking switch according to conventional switching techniques vs. reduce burstiness of notifications in a subscription model) of the two references.

Dependent Claims 18 and 19 inherit the foregoing patentably distinguishing elements of base Claim 17 and should be allowed for at least the same reasons presented above.

Claims 4, 5, and 7 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Hoof in view of Chuang *et al.* (“Matching Output Queueing with a Combined Input/Output-Queued Switch,” hereinafter “Chuang”). Applicants respectfully disagree.

Chuang, which is directed to a combined input/output queued switch, was introduced as a secondary reference against Claim 4 only due to a deficiency of Hoof regarding “determining an incoming cell’s priority based on the time of the cell departing from an output queue and the times of other cells in the output queue to depart” (Office Action, page 9, lines 6-8). Chuang was only introduced against Claim 5 regarding lowest time-to-leave scheduling, lowest time-to-leave blocking, and non-negative slackness insertion. Chuang was only introduced against Claim 7 for emulation of an output queued packet switch. Chuang does not cure the deficiency of Hoof regarding Claim 1, from which Claims 4, 5, and 7 depend, as discussed above (“switch fabric storing cells based on ... output queues”). Therefore, Applicants respectfully submit that the rejection of Claims 4, 5, and 7 under 35 U.S.C. 103(a) is improper and should be withdrawn.

Claim 10 stands rejected under 35 U.S.C. 103(a) as being unpatentable over Hoof in view of Chao and further in view of Chuang. Similar arguments apply for Claim 10 as for Claim 4, discussed above.

Claims 13 and 14 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Hoof in view of Chao and further in view of Rojas-Cessa *et al.* (“CIXB-1: Combined Input-One-Cell Crosspoint Buffered Switch,” hereinafter “Rojas-Cessa”). Applicants respectfully disagree.

Rojas-Cessa, which was introduced as a reference against Claim 13 only for multiple virtual output queues, does not cure the deficiencies of Hoof and Chao regarding Claim 8, from which Claim 13 depends. Therefore, Applicants respectfully submit that the rejection of Claims 13 and 14 (depending from 13) under 35 U.S.C. 103(a) is improper and should be withdrawn.

Claims 15 and 16 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Hoof in view of Chao and further in view of Zhang ("Service Disciplines for Guaranteed Performance Service in Packet-Switching Networks," hereinafter "Zhang"). Applicants respectfully disagree.

Zhang was introduced as a reference against these claims only for FIFO groups (Claim 15) and a plurality of crosspoint schedulers (Claim 16). Zhang does not cure the deficiencies of Hoof and Chao regarding Claim 8, from which Claim 15 depends. Therefore, Applicants respectfully submit that the rejection of Claims 15 and 16 (depending from 15) under 35 U.S.C. 103(a) is improper and should be withdrawn.


CONCLUSION

In view of the above amendments and remarks, it is believed that all pending claims (Claims 1-19) are in condition for allowance, and it is respectfully requested that the application be passed to issue. If the Examiner feels that a telephone conference would expedite prosecution of this case, the Examiner is invited to call the undersigned.

Respectfully submitted,

HAMILTON, BROOK, SMITH & REYNOLDS, P.C.

By


Arvind Jairam

Registration No. 62,759

Telephone: (978) 341-0036

Facsimile: (978) 341-0136

Concord, MA 01742-9133

Date:

7/14/9